

# Optimizing gravitational-wave searches for a population of coalescing binaries: Intrinsic parameters

T. Dent\*

*Albert-Einstein-Institut (Max-Planck-Institut für Gravitationsphysik),  
Callinstr. 38, Hannover, Germany and Leibniz-Universität Hannover, Germany*

J. Veitch†

*Nikhef, Science Park 105, Amsterdam 1098XG, Netherlands*

(Dated: February 21, 2014, LIGO-P1300194-v2)

We revisit the problem of searching for gravitational waves from inspiralling compact binaries in Gaussian coloured noise. If the intrinsic parameters of a quasi-circular, non-precessing binary are known, then the optimal statistic for detecting the dominant mode signal in a single interferometer is given by the well-known two-phase matched filter. However, the matched filter signal-to-noise ratio (SNR) is *not* in general an optimal statistic for an astrophysical population of signals, since their distribution over the intrinsic parameters will almost certainly not mirror that of noise events, which is determined by the (Fisher) information metric. Instead, the optimal statistic for a given astrophysical distribution will be the Bayes factor, which we approximate using the output of a standard template matched filter search. We then quantify the improvement in number of signals detected for various populations of non-spinning binaries: for a distribution of signals uniformly distributed in volume and with component masses distributed uniformly over the range  $1 \leq m_{1,2}/M_{\odot} \leq 24$ ,  $(m_1 + m_2)/M_{\odot} \leq 25$  at fixed expected SNR, we find  $\geq 20\%$  more signals at a false alarm threshold of  $10^{-6}$  Hz in a single detector. The method may easily be generalized to binaries with non-precessing spins.

## I. INTRODUCTION

Although some hints exist to the possible masses and spins of coalescing compact binary (CBC) systems visible to ground-based gravitational-wave (GW) interferometers [1], the dearth of current astronomical observations allows for a very broad uncertainty in the distribution of signals over these “intrinsic” parameters. Thus, searches for GWs from these sources must consider a correspondingly broad parameter space, which leads to increased computational cost and lower search efficiency relative to the case of binaries of known parameter values. Techniques for covering the space in a computationally efficient way are well developed (e.g. [2–4]), but the problem of dealing with the increased number of false alarms arising from a broad search space has received relatively little attention. Given the improved low-frequency sensitivity of Advanced LIGO and Virgo, and the probable need to search for spinning signals (see [5–9] and references therein) implying a yet larger parameter space, this issue is likely to become increasingly urgent.

We build on methods used in previous searches for CBC signals, as described in detail in [10–13]. In these, the intrinsic parameter space is searched by identifying a set (“bank”) of discrete parameter values, each representing a “template” signal, which are sufficiently closely spaced that the possible loss of search sensitivity due to a discrepancy between any true signal and the best fitting template is small and bounded above [14–16]. The matched filter is evaluated for each template, and for each time sample in the experimental data [10, 11]. Then the standard approach for detection of coalescing binaries with unknown intrinsic parameters is (at least in Gaussian noise) to weight all templates equally. This choice is implemented by taking the detection statistic, for candidate events in data from a single interferometric detector, to be the local maximum of a matched filter signal-to-noise ratio (SNR) over the bank, and over coalescence time. The SNR maximum corresponds to the maximum-likelihood value of the signal parameters in Gaussian noise [17]. The resulting “triggers” from each detector are subsequently tested for consistency of parameter values between detectors, and the detection statistic for the resulting coincident events is taken to be the quadrature sum of single-detector SNRs (e.g. [12]); however, we will not consider these multi-detector stages of the analysis here.

In this work we first argue (Section II) that the maximum matched filter SNR is *not* in general an optimum detection statistic in the Neyman-Pearson sense of maximizing detection probability at fixed false alarm probability. This lack of optimality is due to the necessity of marginalizing over the signal parameters when evaluating the Bayes factor (likelihood ratio) for the signal *vs.* noise hypotheses. We show that an approximation of the required integral around a peak of signal likelihood yields, in addition

---

\* thomas.dent@aei.mpg.de

† john.veitch@ligo.org

to a strongly increasing dependence on the SNR, a non-trivial dependence on the intrinsic parameters. This dependence varies as the ratio of the prior probability density of signals to the density of the Fisher information metric for the filters (closely related to the metric used to lay out the template bank [14–16]). Only when these two densities are equal is the matched filter SNR optimal. The use of the maximum SNR as a detection statistic over a broad parameter space thus corresponds to a prior over the binary’s intrinsic parameters which is highly artificial, and may differ considerably from any astrophysically plausible distribution. For instance, the implied prior in a search for non-spinning binary inspirals with component masses 1–24  $M_{\odot}$  and maximum mass 25  $M_{\odot}$  varies by several orders of magnitude over the space of component masses.<sup>1</sup>

Similarly, various detection statistics used in searches for unmodelled gravitational-wave bursts have been shown [18] to correspond to non-physical or implausible priors on the distribution of signals; and the maximization process involved in the “F-statistic” matched filter search for continuous gravitational waves corresponds to a non-physical prior on the source parameters for an individual template [19].

In Section III we discuss implementation in the context of a matched filter search over a standard template bank; we are able to correct for the implicit prior with a negligible increase in computational cost, by computing a likelihood ratio statistic that effectively re-weights template matched filter values via a mass-dependent function incorporating both an assumed astrophysical prior and the Fisher metric density.

Other approaches to using likelihood statistics for coalescing binary signals are described in [20, 21]. However, these studies are primarily concerned with accounting for the presence of non-Gaussian transient noise and/or variable detector sensitivity. The algorithm described in [21], intended primarily to improve search performance in non-Gaussian noise, identifies maxima of a likelihood statistic which is derived from the SNR and a signal-based chi-squared test. This statistic is evaluated over regions of parameter space chosen such that every template is *a priori* weighted equally, thus in Gaussian noise there would still be a one-to-one mapping from SNR to the expected ranking statistic. In contrast, our method accounts for the basic discrepancy between the distributions of signal and noise events that will occur, even in Gaussian stationary noise, across *any* coalescing binary search, and allows us to derive an optimal single-detector statistic for a given signal distribution.

Section IV quantifies how our likelihood-ratio statistic performs on a single-detector search of Gaussian data with low-mass, non-spinning binary inspiral signals; we see improvements of tens of percent in the number of signals recovered at fixed false alarm rate. We also investigate how the performance of the statistic is affected by a discrepancy between the assumed prior over the binary component masses and the true distribution of signals.

Finally in Section V we discuss possible further developments: extension of the re-weighting method to multi-detector searches; possible implementation via machine learning in cases where evaluating the signal and noise densities is complex or computationally costly; relations with likelihood-based methods for optimizing searches in the presence of non-Gaussian noise transients; and implications of the current highly uncertain predictions for coalescing binary source rates [1].

## II. THEORETICAL MOTIVATION

In this section we revisit the standard matched filter statistic for detection of a modelled signal in Gaussian coloured noise. We give a condensed overview of the basic strategy of matched filtering as applied to compact binary searches; readers less familiar with the background may consult previous descriptions [10, 11] and references therein. (We also discuss relevant aspects of implementation in Section III.) We point out that when the signal has one or more unknown parameters, maximization of the matched filter over those parameters effectively introduces a prior on the parameter space which is determined by the information metric density, *i.e.* the Jeffreys prior, or equivalently  $\sqrt{|\det \Gamma^{ij}|}$  where  $\Gamma^{ij}$  is the Fisher matrix. Thus, the matched filter is only an optimal statistic for a distribution of signals which matches this prior. We then find a simple approximation to the optimal detection statistic for a prior distribution which may differ strongly from the Jeffreys one.

### A. Detection of a signal with unknown parameters

Detecting a gravitational wave in data from an interferometric detector can be thought of as a problem of model selection [22]. The aim is to calculate relative probabilities between the two hypotheses, labelled as:

$S$  : There is a signal  $h(f)$  and additive Gaussian noise of power spectrum  $S(f)$  in the data

$N$  : There is no signal, but only noise with power spectrum  $S(f)$  in the data.

---

<sup>1</sup> The density actually vanishes at the equal-mass line  $m_1 = m_2$ , leading to an (integrable) singularity in the implied prior, however even without considering templates near this line we see 4-5 orders of magnitude variation.

For a set of observed data  $d$  we write

$$\frac{P(S|d)}{P(N|d)} = \frac{P(S)}{P(N)} \frac{p(d|S)}{p(d|N)} \equiv \mathcal{B}_{\text{SN}}(d) \frac{P(S)}{P(N)}, \quad (1)$$

where  $\mathcal{B}_{\text{SN}}(d)$  is the likelihood ratio between the hypotheses (also known as the Bayes factor) and  $P(S)/P(N)$  is the prior odds in favour of the signal hypothesis.

For an exactly known signal waveform  $h$  with no parameters, and in Gaussian noise, the likelihood ratio can be written [17, 22, 23]

$$\mathcal{B}_{\text{SN}}(d) = \frac{\exp\left[-\frac{1}{2} \langle d - h | d - h \rangle\right]}{\exp\left[-\frac{1}{2} \langle d | d \rangle\right]} = \exp\left[\langle d | h \rangle - \frac{1}{2} \langle h | h \rangle\right]. \quad (2)$$

Here  $\langle \cdot | \cdot \rangle$  is the usual one-sided noise-weighted inner product defined over the space of possible data streams  $d$ ; for discrete data in the frequency domain with sampling interval  $\Delta f$  we have

$$\langle a | b \rangle \equiv 4\Delta f \operatorname{Re} \left[ \sum_i \frac{a_i^* b_i}{S_i} \right], \quad (3)$$

where  $i$  indexes frequency samples from 0 up to (but not including) the Nyquist frequency. Since  $h$  is constant,  $\mathcal{B}_{\text{SN}}$  is monotonic in  $\langle d | h \rangle$ ; this quantity, called the matched filter, is thus an efficient detection statistic [10, 22]. The Neyman-Pearson lemma states that such a statistic is optimal, yielding the maximum possible detection probability at fixed false alarm probability.

However, in searching for an unknown inspiral signal the waveform arriving at the detector has several free parameters. Thus it is necessary to calculate (an approximation to) the optimal statistic for a signal with unknown parameters and compare this statistic with the standard matched filter. We first consider a parametrised signal without assuming any special form, except concerning the signal amplitude which we take as unknown and continuously variable between 0 and infinity. Write the signal as

$$h(\vec{\theta}) = \mathcal{A} \hat{h}(\vec{\theta}), \quad \text{where} \quad \langle \hat{h} | \hat{h} \rangle = 1, \quad (4)$$

and where  $\vec{\theta}$  is a set of parameter values in the space  $\Theta$ , which we consider as a smooth manifold; the normalized signal  $\hat{h}(\vec{\theta})$  is restricted to a sub-manifold by the condition (4). If considering one specific point in parameter space  $\vec{\theta}$ , the likelihood ratio is

$$L_{\text{SN}}(d; \mathcal{A}, \vec{\theta}) = \frac{p(d|S; \mathcal{A}, \vec{\theta})}{p(d|N)} = \exp\left[\mathcal{A} \langle d | \hat{h}(\vec{\theta}) \rangle - \frac{\mathcal{A}^2}{2} \langle \hat{h}(\vec{\theta}) | \hat{h}(\vec{\theta}) \rangle\right]. \quad (5)$$

Then allowing  $\mathcal{A}$  to vary, we find a maximum of likelihood at  $\mathcal{A}_{\text{ML}} = \langle d | \hat{h}(\vec{\theta}) \rangle = \rho(\vec{\theta})$ , using the standard matched filter notation  $\rho = \langle d | h \rangle / \sqrt{\langle h | h \rangle} = \langle d | \hat{h} \rangle$ , and the likelihood ratio for the specific parameter values chosen is  $p(d|S; \vec{\theta})/p(d|N) = \exp[\rho(\vec{\theta})^2/2]$ .

If we now consider varying the parameters  $\vec{\theta}$ , obviously the point with maximum  $\rho(\vec{\theta})$  also maximizes  $L_{\text{SN}}$ .

This straightforward choice does not, though, solve the problem in hand. The marginalized likelihood ratio including unknown parameters is

$$\mathcal{B}_{\text{SN}}(d) = \frac{\int_{\Theta} d\mathcal{A} d\vec{\theta} p(\mathcal{A}, \vec{\theta}|S) p(d|S; \mathcal{A}, \vec{\theta})}{p(d|N)} = \int_{\Theta} d\mathcal{A} d\vec{\theta} p(\mathcal{A}, \vec{\theta}|S) L_{\text{SN}}(d; \mathcal{A}, \vec{\theta}), \quad (6)$$

where  $p(\mathcal{A}, \vec{\theta}|S)$  is a prior distribution for the signal parameters which could arise from our knowledge of probable astrophysical sources and their evolution. This expression is also an optimal statistic, in the Neyman-Pearson sense, for the signal distribution described by the prior.

In general the integral over  $\Theta$  must be evaluated numerically (*e.g.* using nested sampling [24] or thermodynamic integration [25]), but under some conditions—essentially, the presence of a sufficiently loud signal—it may be approximated via expansion around the maximum which at first non-trivial order yields the “linear signal approximation” (LSA: see for instance [26, 27]). This approximation assumes a well-defined peak of the likelihood, Eq. (5), for a signal of parameters  $\vec{\theta}^{\text{ML}}$  within  $\Theta$ , near which the inner products involving  $h(\vec{\theta})$  vary smoothly over the parameter space. [28] finds good agreement between this approximation and a full Bayes factor calculation in the high SNR limit when analysing white dwarf binary systems in mock LISA data. Due to the exponential dependence of  $L_{\text{SN}}$  on the matched filter  $\rho(\vec{\theta})$  (always assuming Gaussian noise) the integral over the parameter space is likely to be dominated by a domain near the maximum likelihood point; furthermore, although secondary (disconnected) maxima of likelihood will generally occur, their contributions to the Bayes factor will in almost all cases

be subdominant compared to the global maximum. Here we approximate the result by the integral over a single peak; summing the evidence over possible multiple peaks identified by the search would be a refinement to this approximation.

We make the further approximation that the prior  $p(\mathcal{A}, \vec{\theta}|S)$  is sufficiently slowly-varying near the peak that it may be taken as constant over regions of  $\Theta$  that contribute significantly to the integral; this allows us to take the prior value at the peak  $p(\mathcal{A} = \mathcal{A}_{\text{ML}}, \vec{\theta} = \vec{\theta}^{\text{ML}}|S)$  outside the integral.<sup>2</sup> The usefulness of this approximation will clearly depend on the actual behaviour of the prior over such regions, whose size and shape are determined by the variation of the likelihood. Since the likelihood is increasingly strongly peaked for higher-amplitude signals, the constant prior approximation will become increasingly good in the strong signal limit, which is also where the LSA is valid.

If this, or other of our approximations, fail badly enough, the ranking statistic we derive will not be (close to) optimal for the signal distribution considered, and may not represent an improvement in detection efficiency over the maximum matched filter  $\rho_{\text{max}}$ . We will see this is the case over a relatively small subset of the parameter space considered in our example application (Section IV): our chosen prior, while uniform in component masses, has a singularity at the line  $m_1 = m_2$  in coordinates for which the likelihood is well described by a Gaussian peak of predictable width. Naïvely applying the constant prior approximation for maxima sufficiently close to this singularity could produce an error of indefinitely large size, which could result in random noise fluctuations being wrongly assigned a high ranking value, hurting the efficiency of the statistic. In principle we could go beyond this approximation by expanding the prior about the maximum-likelihood point in appropriate coordinates, or by attempting to explicitly carry out the integral (6). However this would go beyond the intention of our study, which is to quantify the *leading-order* effects of the signal prior; also, in practice we find that a simple ad-hoc regularization of the singularity leads to a statistic which is demonstrably more efficient than  $\rho_{\text{max}}$ , and we find no evidence of any substantial loss in efficiency due to our simple modelling of the prior: see Section IV B.

The maximum likelihood condition implies

$$\begin{aligned} \mathcal{A}_{\text{ML}} &= \langle d | \hat{h}^{\text{ML}} \rangle \equiv \rho_{\text{max}}, \\ \langle \partial_i \hat{h}^{\text{ML}} | d - \hat{h}^{\text{ML}} \rangle &\equiv \left\langle \frac{\partial \hat{h}}{\partial \theta_i} \middle| d - \hat{h}^{\text{ML}} \right\rangle_{|\vec{\theta}^{\text{ML}}} = 0, \end{aligned} \quad (7)$$

and the normalization of the filters  $\langle \hat{h}(\vec{\theta}) | \hat{h}(\vec{\theta}) \rangle = 1$  implies further that

$$\langle \hat{h}(\vec{\theta}) | \partial_i \hat{h}(\vec{\theta}) \rangle = 0, \quad (8)$$

$$\langle d | \partial_i \hat{h}^{\text{ML}} \rangle = \left( \frac{\partial \rho}{\partial \theta_i} \right)_{|\vec{\theta}^{\text{ML}}} = 0. \quad (9)$$

We expand about the maximum to first derivative order in the waveform  $\partial_i \hat{h}$ , where now the parameters indexed by  $i$  will not include the amplitude, but keep terms to *second* order in parameter deviations  $\Delta\theta_i = (\theta_i - \theta_i^{\text{ML}})$ .<sup>3</sup> We enforce the norm condition to this order via

$$\begin{aligned} \hat{h}(\vec{\theta}^{\text{ML}} + \Delta\vec{\theta}) &= C \left[ \hat{h}^{\text{ML}} + \Delta\theta_i \partial_i \hat{h}^{\text{ML}} + \dots \right], \\ C^2 \left[ \langle \hat{h}^{\text{ML}} | \hat{h}^{\text{ML}} \rangle + 2\Delta\theta_i \langle \hat{h}^{\text{ML}} | \partial_i \hat{h}^{\text{ML}} \rangle + \Delta\theta_i \Delta\theta_j \langle \partial_i \hat{h}^{\text{ML}} | \partial_j \hat{h}^{\text{ML}} \rangle \right] &= 1 \\ \implies C &\simeq 1 - \frac{1}{2} \Delta\theta_i \Delta\theta_j \langle \partial_i \hat{h}^{\text{ML}} | \partial_j \hat{h}^{\text{ML}} \rangle \equiv 1 - \frac{1}{2} \tilde{g}_{ij}(\vec{\theta}^{\text{ML}}) \Delta\theta_i \Delta\theta_j, \end{aligned} \quad (10)$$

where we have pre-emptively defined the parameter space metric  $\tilde{g}_{ij} = \langle \partial_i \hat{h} | \partial_j \hat{h} \rangle$  (see [14, 15]). For the amplitude we write  $\mathcal{A} = \mathcal{A}_{\text{ML}} + \Delta\mathcal{A}$ . Expanding inside the exponential in Eq. (5) around the maximum likelihood point and using the relations (7-9) we obtain

$$\log L_{\text{SN}}(d; \mathcal{A}, \vec{\theta}) \simeq C \mathcal{A}_{\text{ML}} (\mathcal{A}_{\text{ML}} + \Delta\mathcal{A}) - \frac{(\mathcal{A}_{\text{ML}} + \Delta\mathcal{A})^2}{2} \simeq \frac{1}{2} \left[ \mathcal{A}_{\text{ML}}^2 (1 - \tilde{g}_{ij}(\vec{\theta}^{\text{ML}}) \Delta\theta_i \Delta\theta_j) - \Delta\mathcal{A}^2 + \dots \right]. \quad (11)$$

Thus the metric  $\tilde{g}_{ij}$  describes the fall-off in log likelihood as a function of the deviation of the parameters from their ML values; it is proportional to the Fisher information matrix (at fixed SNR). The *match*  $M$ , *i.e.* the overlap of normalized waveforms, used in constructing template banks for coalescing binary searches [15], is also described by this metric via

$$M(\vec{\theta}_*, \Delta\vec{\theta}) = \langle \hat{h}(\vec{\theta}_*) | \hat{h}(\vec{\theta}_* + \Delta\vec{\theta}) \rangle \simeq 1 - \frac{1}{2} \tilde{g}_{ij}(\vec{\theta}_*) \Delta\theta_i \Delta\theta_j, \quad (12)$$

<sup>2</sup> Under this approximation it is also the case that the posterior distribution  $p(\mathcal{A}, \vec{\theta}|S)p(d|S; \mathcal{A}, \vec{\theta})$  varies as the signal likelihood near the maximum.

<sup>3</sup> Second derivatives of  $\hat{h}$  will also contribute to this order in  $\Delta\theta_i$ , but at sub-leading order in amplitude  $\mathcal{A}_{\text{ML}}$  (see *e.g.* [26]), thus we omit them here.

for a small deviation  $\Delta\vec{\theta}$  away from  $\vec{\theta}_*$ .

The integral in Eq. (6) may be performed approximately to yield

$$\mathcal{B}_{\text{SN}}(d) \simeq p(\mathcal{A} = \rho_{\text{max}}, \vec{\theta} = \vec{\theta}^{\text{ML}} | S) \exp\left(\frac{\rho_{\text{max}}^2}{2}\right) \rho_{\text{max}}^{-m} \sqrt{\frac{(2\pi)^{m+1}}{\det \tilde{g}(\vec{\theta}^{\text{ML}})}}, \quad (13)$$

where  $m$  is the dimensionality of the parameter space *excluding* the amplitude. This expression is no longer monotonic in the maximum matched filter value  $\rho_{\text{max}}$  (for  $\rho_{\text{max}} > 1$ ).<sup>4</sup> Including the leading dependence of  $\mathcal{B}_{\text{SN}}$  on the parameters of the maximum has two effects: firstly via the metric density  $\sqrt{\det \tilde{g}}$ , and secondly by allowing a prior over the parameter space. Note that the ratio of the prior density to the metric density is invariant under a change of coordinates over  $\Theta$ .

The approximate Bayes factor of Eq. (13) is only monotonic in  $\rho_{\text{max}}$  if the signal prior is numerically equal to the metric density. Thus, *maximizing the matched filter over a parameter space implicitly introduces a prior over the space: the maximum matched filter is not an optimal statistic unless the signal distribution is equal to the metric density*. The implicit prior is given by  $p_{\text{imp}}(\mathcal{A}, \theta | S) = f(\mathcal{A}) \sqrt{\det \tilde{g}(\theta)}$ , where  $f(\mathcal{A})$  is a smooth function of signal amplitude such that  $\mathcal{B}_{\text{SN}}$  is a monotonically increasing function of  $\rho_{\text{max}}$  at  $\mathcal{A} = \rho_{\text{max}}$  ( $\rho_{\text{max}} > 1$ ).

A similar implicit prior has been investigated in the case of amplitude parameters of continuous gravitational waves from isolated neutron stars, where it was indeed found that accounting for the correct prior could lead to improvements in detection efficiency over the maximized matched-filter (“F-statistic”) method [19]. The use of explicit, physically motivated priors has also been shown to improve the efficiency of searches for unmodelled (or, to be more accurate, weakly-modelled) GW bursts [18] compared to standard analysis methods. Note that an optimisation of template bank placement using a prior signal distribution was proposed in [29], but under a criterion of maximum efficiency at fixed computing cost (taken as proportional to the number of templates) rather than at fixed false alarm rate.

## B. Search parameters of coalescing binary signals

We now explore the implications of this result for binary coalescence signals by considering their parameters as seen, for simplicity, by a single interferometric detector. (Multi-detector searches will be discussed in section V as a topic of future development.) The gravitational wave signal from a binary inspiral is dependent on up to 15 parameters, in the case of a binary with a quasi-circular orbit and spinning components (neglecting matter effects). We write the signal in the form  $h(f) = \mathcal{A}\hat{h}(f; \vec{M}) \exp(-2\pi i f t_c - i\phi_c)$ , where the “intrinsic parameters”  $\vec{M}$  contain the masses and the spin vectors of the binary components, separating out the dependence on the coalescence time of the signal  $t_c$  and coalescence phase  $\phi_c$ . For dominant-mode signals in ground-based detectors for which the detector antenna pattern is approximately constant over their duration, the phase evolution  $\hat{h}(f; \vec{M})$  is a function of intrinsic parameters only, whereas the amplitude  $\mathcal{A}(\vec{M}, \alpha, \delta, \psi, \iota, d_L)$  may depend on the intrinsic parameters and on the inclination  $\iota$  angle of the binary, the polarisation angle  $\psi$ , the luminosity distance  $d_L$ , and the position on the sky relative to the detector  $\alpha, \delta$ . Together, these parametrise the signal hypothesis  $\{\vec{M}, \mathcal{A}, t_c, \phi_c\} \in \Theta$ .

## C. Coalescence phase and the coincidence metric

As described in [10, 11], in searching for binaries where the orbital plane does not precess, the coalescence phase  $\phi_c$  is not explicitly searched over in matched filtering single-detector data. Rather, a complex matched filter is evaluated via Fast Fourier Transform and the analogue of the matched filter for known signals is found to be the absolute value of the complex filter. For a single filter template (a point in parameter space  $\vec{M}$ ) the likelihood ratio marginalized over  $\phi_c$  with uniform prior was shown to be monotonic in the “two-phase” SNR  $\rho = |z|/\sigma$ , where  $z = \langle d|h(\vec{M}, t_c, \phi_c = 0) \rangle + i \langle d|h(\vec{M}, t_c, \phi_c = \pi/2) \rangle$  and  $\sigma$  is a template normalization factor;  $\rho$  is thus an optimal statistic for a signal of unknown phase. The maximum  $\rho$  value is also equal to the maximum-likelihood amplitude  $\mathcal{A}_{\text{ML}}$ .

Since the dependence of the signal on  $\phi_c$  is simple and exactly known, we may similarly marginalize the likelihood ratio over  $\phi_c$  before expanding in the amplitude and other parameters; for a detailed calculation refer to appendix A. This is expected to be more accurate than a quadratic expansion around the maximum likelihood value; we explicitly show, though, that the two methods agree for high-amplitude signals. The phase-marginalized likelihood ratio is then

$$L'_{\text{SN}}(d; \mathcal{A}, \vec{\theta}) \equiv \int \frac{d\phi_c}{2\pi} L_{\text{SN}}(d; \phi_c, \mathcal{A}, \vec{\theta}) \simeq \exp\left(-\frac{\mathcal{A}^2}{2}\right) I_0 \left[ \mathcal{A} \sqrt{\langle \hat{h}_0(\vec{\theta}) | d \rangle^2 + \langle \hat{h}_{\pi/2}(\vec{\theta}) | d \rangle^2} \right],$$

<sup>4</sup> Although this approximate formula for the Bayes factor grows without limit as  $\rho_{\text{max}} \rightarrow 0$ , this does *not* indicate an increasing preference for signal over noise in the low-SNR limit. First, the approximation is only valid for (at least moderately) high SNR, certainly not for  $\rho \lesssim 1$ ; second, in any case the probability that the *maximum* matched filter value over any extended region of parameter space is  $\lesssim 1$  is vanishingly small.

where  $I_0$  is the modified Bessel function of the first kind; this may be expanded in the other parameters and marginalized to yield

$$\int d\mathcal{A} d\vec{\theta} p(\mathcal{A}, \vec{\theta}|S) L' \simeq p(\mathcal{A} = \mathcal{A}_{\text{ML}}, \vec{\theta} = \vec{\theta}^{\text{ML}}|S) L'_{\text{ML}} \int d(\Delta\mathcal{A}) e^{-\Delta\mathcal{A}^2/2} \int d^N \vec{\theta} \exp \left[ -\frac{\mathcal{A}_{\text{ML}}^2}{2} \gamma_{pq} \Delta\theta_p \Delta\theta_q \right] \quad (14)$$

$$\simeq p(\mathcal{A} = \rho_{\text{max}}, \vec{\theta} = \vec{\theta}^{\text{ML}}|S) \exp \left( -\frac{\rho_{\text{max}}^2}{2} \right) I_0(\rho_{\text{max}}^2) \rho_{\text{max}}^{-n} \sqrt{\frac{(2\pi)^{n+1}}{\det \gamma(\vec{\theta}^{\text{ML}})}}, \quad (15)$$

where as before we identify  $\mathcal{A}_{\text{ML}}$  with the local maximum matched filter  $\rho_{\text{max}}$ ;  $p, q$  label the  $n$  parameters other than  $\phi_c$  and  $\mathcal{A}$ ; and  $\gamma_{pq}$  is the phase-projected metric

$$\gamma_{pq} = \tilde{g}_{pq} - \frac{\tilde{g}_{p\phi_c} \tilde{g}_{q\phi_c}}{\tilde{g}_{\phi_c\phi_c}} = \left[ \langle \partial_p \hat{h}_0 | \partial_q \hat{h}_0 \rangle - \langle \partial_p \hat{h}_0 | \hat{h}_{\pi/2} \rangle \langle \partial_q \hat{h}_0 | \hat{h}_{\pi/2} \rangle \right]. \quad (16)$$

This metric was used in [30] to describe the dependence of the phase-maximized likelihood on the remaining signal parameters, *i.e.* coalescence time and component masses. Its inverse gives the expected error matrix over these parameters for the maximum-likelihood point, compared to the true values when a signal is present; it was therefore used to implement a consistency test for “triggers” recorded at maxima of the matched filter  $\rho$  in the data streams from different detectors [30].

#### D. Prior parameter distribution

In order to evaluate an optimized statistic we need to specify the prior  $p(\mathcal{A}, \vec{M}, t_c, \phi_c|S)$ . The coalescence phase and time dimensions are trivial and have uniform priors; also, at least for stationary detector noise, the distributions over other parameters do not depend on the values of  $t_c$  and  $\phi_c$ .<sup>5</sup> Further, for a distribution of coalescing binary sources uniform in physical volume, and where there are no correlations between the luminosity distance and any other physical parameter, one may show (see for instance [31]) that there is a universal distribution of signals over the amplitude  $\mathcal{A}$ ,  $p(\mathcal{A}|S) \propto \mathcal{A}^{-4}$ , thus the prior may be factorized as

$$p(\mathcal{A}, \vec{M}, t_c, \phi_c|S) \propto p(\mathcal{A}|S) p(\vec{M}|S, \mathcal{A}) \propto \mathcal{A}^{-4} p(\vec{M}|S, \mathcal{A}), \quad (17)$$

where  $p(\vec{M}|S, \mathcal{A})$  is the distribution over the intrinsic parameters spanning the template bank at constant expected SNR. We will discuss this prior distribution further in Section III C where it will appear as the rate distribution of signals of a given SNR over the search parameter space.

### III. IMPLEMENTATION IN MATCHED FILTER SEARCH

In this section we describe some currently used methods for searching a parameter space of coalescing binary systems and show how the output of such a search corresponds to the statistical derivation of the previous section. We are thus able to implement our approximately optimal statistic using the “triggers” produced by standard search algorithms [11] in combination with a parameter-dependent re-weighting.

#### A. Template bank layout

To begin, the template bank, a set of discrete parameter values  $\vec{M}_k$ , is chosen in order to minimize the maximum possible “mismatch” between a signal waveform and the best fitting template. The complex matched filter is evaluated for each of the templates, and maxima of its modulus  $\rho$  over the bank, and over short time intervals, are recorded. For dominant-mode signals from non-precessing binary systems, the amplitude, coalescence time and coalescence phase can be efficiently searched by a single Fourier transform, leaving the two binary masses and possible non-precessing spin components as dimensions to be covered by the bank; we will refer to these as “intrinsic parameters”.

Each template “covers” a finite region of the intrinsic parameter space, *i.e.* has a sufficiently high match  $M$  with signals in the region. The match, for a signal with intrinsic parameters  $\vec{M}_S$ , is the ratio of the expected SNR recovered by a template at

<sup>5</sup> However, if the detector sensitivity changes with time the prior distribution over  $t_c$  of signals with a given expected SNR, *i.e.* a given value of  $\mathcal{A}$ , is *not* uniform. We will not consider that case here.

$\vec{M}_S + \delta\vec{M}$ , after maximizing over time and phase, to the optimal SNR recovered by a filter with the signal's parameters  $\vec{M}_S$ . For broad-band detectors, the total region where astrophysical signals may be present is much larger than that covered by a single template, therefore we require a “bank” of many templates; these are typically laid out in a grid designed to achieve a specific minimal match  $M_{\min}$  with at least one template for any signal in the total region (in recent searches  $M_{\min} = 0.97$ ) [2, 3, 15, 16].

The template bank is created with the help of a  $N$ -dimensional metric  $g(\vec{M})$  over the intrinsic parameter space,<sup>6</sup> which is simply a projection along  $t_c$  of the coincidence test metric:

$$g_{kl} = \gamma_{kl} - \frac{\gamma_{tck}\gamma_{tcl}}{\gamma_{tctc}}. \quad (18)$$

As a result of this construction, the density of an ideal template bank over the parameter space is proportional to the metric density [15, 16] (in some coordinates for  $\vec{M}$ ):

$$n_T(\vec{M}) = \text{const.} \times (1 - M_{\min})^{-N/2} \sqrt{\det g(\vec{M})}. \quad (19)$$

In practice, geometrically constructed banks (as described in [2, 3]) for non-spinning binaries may make use of the near-constancy of metric components in coordinates  $\tau_0, \tau_3$  defined as

$$\tau_0 = \frac{A_0}{\eta} (m_1 + m_2)^{-5/3}, \quad \tau_3 = \frac{A_3}{\eta} (m_1 + m_2)^{-2/3}, \quad (20)$$

where  $\eta \equiv m_1 m_2 / (m_1 + m_2)^2$  and  $A_0, A_3$  are numerical constants. Then a regular lattice with fixed spacing and orientation over  $\tau_0, \tau_3$ , implicitly using a metric that is constant and flat over this space, yields a bank for which the deviation from the theoretically correct, fixed minimal match  $M_{\min}$  over the whole parameter space is relatively small for low-mass (non-spinning) binary systems [4]. The hexagonal bank placement used for recent searches of LIGO and Virgo data (see *e.g.* [12, 13, 32]) uses a refinement of the regular lattice via an adaptive algorithm that locally adjusts the grid spacing and orientation to account for gradual variation of the metric components over  $\tau_0, \tau_3$  space [3].

## B. Matched filter triggers and clustering

Each template  $\vec{M}_k$  is used to filter the data stream from a detector  $d(t)$  producing a time series  $\rho_k(t_c)$  of the complex matched filter modulus. In Gaussian noise, all template time series have the same distribution [10],  $p(\rho_k(t_c)) = \rho_k e^{-\rho_k^2/2}$ . These matched filter values show strong correlations over nearby coalescence times and between nearby templates; the filter correlations over small changes in  $t_c$  and  $\vec{M}$  can be described by the  $(N + 1)$ -dimensional “coincidence” metric  $\gamma(t_c, \vec{M})$ .

We would like to obtain an optimal statistic to distinguish noise events from signals, where events are considered to be localized in time and over template intrinsic parameters. In order to use the familiar statistics of false alarms, false dismissals, etc., we will model candidate signals and false alarms by time-localized events (approximately) obeying Poisson statistics. A process of *clustering* relates the set of matched filter time series to such localized events (“triggers”).

*Clustering over time* A signal is expected to give rise to a maximum of the matched filter and also to large values in neighbouring time samples, dying away at other times according to the template’s autocorrelation. We wish to allow for multiple signals (or noise events) over a search time, therefore we record local maxima of the matched filter at intervals shorter than the expected waiting time between signals. These maxima are called *triggers* and will be notated<sup>7</sup> as  $\{t_c, \rho, \vec{M}_k\}$ . In practice, the maximization window is set to 1 s or longer, up to the length of the template, in order to reduce computational storage requirements, and triggers are also not recorded if  $\rho(t_c)$  does not exceed a predetermined threshold, typically  $\rho_{\text{th}} = 5.5$ . For a single template, triggers may be modelled as Poisson events with a  $\rho$ -dependent rate  $\lambda_1(\rho)$ , up to the condition that no more than one trigger may occur within a short time window; the approximation becomes increasingly good for smaller  $\lambda_1$  (and thus increasingly large  $\rho$ ).

*Clustering over the template bank* Any given high-SNR signal or loud noise fluctuation will still produce strongly-correlated maxima above threshold in nearby templates and at similar times. We may again consider recording only the maximum value of  $\rho(\vec{M}_k)$  over a correlated region of template parameters and time; however, this procedure implicitly assumes a uniform prior signal distribution *over different templates*. In other words, maximizing  $\rho$  over templates can only yield an optimal statistic if we are equally likely to see a signal in each template during any given stretch of data.

Since we intend to investigate the role played by the signal distribution in the search, if we were to cluster by maximizing  $\rho$  over the entire bank and over substantial periods of time we might obliterate the effect. Nevertheless, over *small* regions of

<sup>6</sup> Not to be confused with the general parameter metric  $\tilde{g}$  considered earlier!

<sup>7</sup> The matched filter value  $\rho$  should be understood as a trigger value rather than a time series unless a dependence on time  $\rho(t_c)$  is explicitly given.

parameter space encompassing, say, order(10) templates (a typical full bank contains many thousands) we may consider the prior probability of signal as approximately constant between templates, and certainly as varying much less rapidly than the likelihood ratio  $L_{SN} \sim \exp(\rho^2/2)$ . Hence we use the “trigscan” algorithm [33] to cluster by recording maxima of  $\rho$  locally over template parameters and time, proximity being tested via the  $(N + 1)$ -dimensional overlap metric  $\gamma(t_c, \vec{M})$ . Triggers widely separated in parameter space (whether occurring at the same time or different times) will not affect one another through clustering and can be treated as independent.

Note also that if the bank has regions where the density of templates is larger than the theoretical value required for the desired minimal match, i.e. areas of over-coverage, for example due to edge effects, then the number of *unclustered* triggers with a given  $\rho$  will also show an excess in these regions; however, the behaviour of triggers *after* clustering should not be affected by such over-coverage.

### C. Trigger statistics and an optimized event ranking

Given a set of clustered triggers, the problem remaining is to find a detection statistic  $\Lambda(t_c, \rho, \vec{M})$  which will maximize the efficiency of detection at a given false alarm rate. Triggers will be ranked by  $\Lambda$  value and those above a constant threshold  $\Lambda_t$  will be considered as detection candidates. The false alarm rate is the total number of triggers that exceed  $\Lambda_t$ , divided by the analysis time, for data containing no signals; the detection efficiency, for a finite population of signals with a given distribution over physical parameters, is the fraction of signals for which an associated trigger can be identified whose statistic exceeds  $\Lambda_t$ . A non-trivial dependence on  $t_c$  will be appropriate if the detector sensitivity varies with time or if some times are more likely to yield signals than others; here we consider a case where the prior over  $t_c$  is constant, thus  $\Lambda = \Lambda(\rho, \vec{M})$ .

Note that the ranking is determined by comparing statistic values calculated for *single* events or triggers, a choice which corresponds to the procedure of approximating the Bayes factor by an integral over a single maximum of likelihood in the previous derivations.<sup>8</sup>

The Neyman-Pearson optimal ranking is then [20]

$$\Lambda_{\text{opt}}(\rho, \vec{M}) = \frac{\lambda_S(\rho, \vec{M})}{\lambda_N(\rho, \vec{M})} \quad (21)$$

where  $\lambda_S, \lambda_N$  are Poisson rate densities of clustered signal and noise triggers over  $\rho$  and over the template parameters. To relate this to the familiar ratio of likelihoods, consider a short period of data over which the expected number of signal or noise triggers is small ( $\ll 1$ ), and suppose we find one trigger in that time with parameters  $\rho, \vec{M}_k$ : the likelihood ratio  $p(1; \rho, \vec{M}_k|S)/p(1; \rho, \vec{M}_k|N)$ , which is an optimal statistic for any one trigger, is then identical to Eq. (21).<sup>9</sup> Although we may not be able to determine the absolute rates of clustered triggers  $\lambda_{S,N}(\rho, \vec{M})$  theoretically or from prior information, we can constrain their functional forms.

*Modelling noise triggers* In Gaussian noise, each template produces a predictable distribution of matched filter time samples  $p(\rho_k(t_c)) = \rho_k e^{-\rho_k^2/2}$ . However, we are concerned with the distribution of *independent* (clustered) noise events. The clustering procedure will alter the distribution of trigger  $\rho$  values to a function  $p(\rho|N)$ , which we will model empirically. The *rate* of noise triggers with given  $\rho$  may also vary over the parameter space. Since the correlations between triggers are described by the metric  $\gamma(t_c, \vec{M})$ , which is also used to perform clustering, we expect that the noise trigger rate density will be given by the metric density in some coordinates over  $\vec{M}$ , thus:

$$\lambda_N(\rho, \vec{M}) \propto p(\rho|N) \sqrt{\det \gamma(t_c, \vec{M})}. \quad (22)$$

We note in passing that the variation of the metric density  $\sqrt{\det \gamma}$  over the 2-dimensional space of non-spinning component masses  $m_1, m_2$ , for the “low-mass” space  $1 \leq m_{1,2} \leq 24$   $m_1 + m_2 \leq 25$ , is similar to that of the bank metric  $\sqrt{\det g(m_1, m_2)}$ . Thus we expect that regions containing equal numbers of templates will yield approximately equal numbers of clustered noise triggers; we were able to check this explicitly for an example search (see Section IV). Small variations in the per-template rate of clustered noise triggers may be due to the variation of the  $t_c$ -components of  $\gamma$ , representing the filter autocorrelations [4].

*Modelling triggers from an astrophysical signal population* The rate of signal triggers expected depends on both the astrophysical coalescence rates and a number of other factors that determine the analysis sensitivity at given SNR (see for instance [34]). We start with an astrophysical rate density  $R(\vec{M}_S)$ , a function of the source intrinsic parameters  $\vec{M}_S$ , measured in coalescences per year, per cubic Mpc, per volume of parameter space  $d^N \vec{M}_S$ . We assume that coalescences are uniformly distributed

<sup>8</sup> To go beyond this assumption one would evaluate a detection statistic over sets containing multiple triggers, with each set representing the result of a search over an extended data set.

<sup>9</sup> The probabilities for a single template  $\vec{M}_k$  are related to the densities  $\lambda_{S,N}(\vec{M})$  over the bank by a change of coordinates which cancels in this ratio.



over time and space out to the limits of the detector’s maximum sensitive distance, and neglect any effects of redshift. The expected SNR of a signal  $\bar{\rho}$  is affected by the source’s intrinsic parameters (mass and spin), angular extrinsic parameters (orientation and sky direction relative to the detector), the distance to the detector, and the detector’s sensitivity, i.e. its noise power spectral density (PSD).

For a non-spinning inspiral signal where the template’s upper frequency cut-off is above the detector sensitive band, we have  $\bar{\rho} \propto \mathcal{M}^{5/6}/D$ , where  $D$  is the distance to the source and the chirp mass is defined as  $\mathcal{M} \equiv (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ . The proportionality depends on angular parameters and on the detector sensitivity, but this dependence is universal and thus factors out of the relative rates of signals with different  $(\bar{\rho}, \vec{M}_S)$ .

Formally, considering events within an element of the space of angular extrinsic parameters, but explicitly allowing for a varying distance  $D$ , we have a number density of signals at the detector

$$dN \equiv \mu_S(\bar{\rho}, \vec{M}_S) dt d\bar{\rho} d^N \vec{M}_S = \text{const.} \times R(\vec{M}_S) dt dV(\bar{\rho}, \vec{M}_S) d^N \vec{M}_S \quad (23)$$

where  $\mu_S$  is a rate density over the expected SNR  $\bar{\rho}$  and source intrinsic parameters  $\vec{M}_S$ , and  $V(\bar{\rho}, \vec{M}_S)$  is the volume of space over which a signal with the given parameters would have an expected SNR greater than  $\bar{\rho}$ . Since  $V = \text{const.} \times D^3 = \text{const.} \times \mathcal{M}^{5/2} \bar{\rho}^{-3}$ , we find

$$\mu_S(\bar{\rho}, \vec{M}_S) = \text{const.} \times R(\vec{M}_S) \frac{\mathcal{M}_S^{5/2}}{\bar{\rho}^4}. \quad (24)$$

This relation holds with some constant of proportionality for any value of angular parameters, therefore it will also hold for the distribution as a whole provided that the distribution over angular parameters does not depend on  $D$  or  $\vec{M}_S$ .

For signals where the expected SNR does not scale with the chirp mass as above (e.g. when late inspiral or merger lies within the detector’s sensitive band),  $\mathcal{M}_S^{5/2}$  may be replaced by  $D_{\text{hor}}(\vec{M}_S)^3$ , where  $D_{\text{hor}}$  is the horizon distance at which a coalescence with optimal orientation overhead from the detector would have a given, fixed SNR (usually taken as  $\bar{\rho} = 8$ ).

Now consider the distribution of signal triggers  $\lambda_S(\rho, \vec{M})$  over the *recovered* maximum matched filter  $\rho$  and parameter values  $\vec{M}$ . For a given signal  $\rho$  will not be equal to the expected SNR, nor will the ML parameters be equal to those of the source. This parameter inaccuracy is in general caused both by random noise, and by bias in the template waveform model which may cause the best match to occur at recovered parameter values substantially different from the true physical ones; see [35, 36] for relevant discussions of PN inspiral waveform systematics.

In general one should perform a Monte Carlo study to determine the relation between  $\mu_S$  and  $\lambda_S$ , but we note that, for “faithful” templates with relatively small parameter bias, if the signal distribution  $\mu_S$  has sufficiently broad support and is slowly-varying, then the difference between the two distributions is unlikely to be substantial except at low  $\bar{\rho}$  where the recovered trigger distribution will obviously be cut off by the threshold imposed in the analysis. Hence we may take

$$\lambda_S(\rho, \vec{M}) \begin{cases} \simeq \mu_S(\bar{\rho} = \rho, \vec{M}_S = \vec{M}), & \bar{\rho} > \rho_{\text{th}} \\ = 0, & \bar{\rho} < \rho_{\text{th}}. \end{cases} \quad (25)$$

If  $\mu_S(\bar{\rho}, \vec{M}_S)$  does vary strongly over its arguments, for example being strongly peaked over mass or with sharp cutoffs inside the search parameter space, the errors in recovered parameters must be treated more rigorously. The “slowly-varying” condition for Eq. (25) to be valid is analogous to that for neglecting the variation of the prior density when marginalizing over a peak of likelihood.

We then find the approximately optimized statistic, for the source population described by  $R(\vec{M}_S)$ :

$$\Lambda_{\text{opt}}(\rho, \vec{M}_k) \simeq \frac{1}{\rho^4 p(\rho|N)} \frac{R(\vec{M}_S = \vec{M}_k) D_{\text{hor}}(\vec{M}_k)^3}{\sqrt{\det \gamma_{|\vec{M}_k}}}, \quad (26)$$

recalling that  $\vec{M}_k$  is the parameters of the  $k$ ’th template.

Although this statistic is the ratio of rate densities over the parameters of clustered triggers, whereas the Bayes factor evaluated approximately in Eq. (13), (15) is a ratio of marginalized likelihoods for the detector time series, the two are identical in form up to a constant numerical factor if we identify the prior signal distribution with  $R(\vec{M}) D_{\text{hor}}^3 / \rho^4$  and the remaining factors with the inverse of the noise trigger density. This indicates we are solving essentially the same optimization problem in both cases given our approximations and assumptions; specifically, that we consider sufficiently high maxima of likelihood (at least above the threshold for trigger generation), and we marginalize or cluster triggers over regions of parameter space containing no more than one such maximum.




FIG. 1. Metric density calculated by the standard LAL template bank code at each  $(m_1, m_2)$  point in the “low-mass” bank, transformed to  $(t_c, m_1, m_2)$  coordinates, and regularized as described in Section IV B by adding a small constant (equal to 0.3, in the units shown) in quadrature. The apparent excess density of templates near  $m_1 = m_2$  is necessary to avoid potential under-coverage at high  $\eta$ .

#### D. Comparison with SNR ranking

The optimized statistic for a given signal population in Gaussian noise differs in three respects from a ranking of events on matched filter SNR  $\rho$ . First, via the astrophysical rate distribution; second, via the dependence of sensitive volume on the signal parameters; third, via the expected rate density of noise triggers. If, as would be the standard choice in Gaussian noise, we were to rank events via  $\rho$  alone, our search would still be optimized for *some* signal population, specifically that described by

$$R_\rho(\vec{M}_S) = \sqrt{\det \gamma_{|\vec{M}_S}} \cdot D_{\text{hor}}(\vec{M}_S)^{-3}.$$

It may be plausible that the astrophysical coalescence rate varies as some negative power of the horizon distance, at least within some restricted space of masses, but there is no physical reason why we would expect it to vary proportionally to the metric density associated with the search filters, which may depend on both the detector noise PSD and the signal model used. Furthermore, many different searches may be performed which will generally have different values of  $\gamma(t_c, \vec{M})$ , whereas the astrophysical population will have a fixed and constant distribution  $R(\vec{M}_S)$ .

Still, if the variation in  $\sqrt{\det \gamma}$  over the search space were negligibly small compared to variations in likelihood due to differences in  $\rho$ , then the matched filter value might still be a good approximation to an optimal statistic. We show, though, that this is *not* the case for the standard “low-mass” inspiral search space in Fig. 1, where the metric density evaluated by a standard LIGO Algorithm Library (LAL) [37] bank placement code at 2PN order, transformed to  $t_c, m_1, m_2$  coordinates, is plotted for templates with component masses between 1–24  $M_\odot$  and total mass 2–25  $M_\odot$ , for a lower frequency cutoff of 40 Hz and a noise PSD representative of the “early advanced LIGO” sensitivity range given in [38]. There is a variation of several orders of magnitude between the sparsest region and the densest near  $m_1 = 1.5, m_2 = 1$ . Hence, we may expect a significant difference between the ranking of events by an optimized statistic  $\Lambda_{\text{opt}}$  compared to ranking by  $\rho$ .

### IV. EXAMPLE APPLICATION: SEARCH FOR LOW-MASS NON-SPINNING BINARY COALESCENCE

In this section we will quantify the differences in search efficiency produced by using the optimized statistic described in the previous section for a simple example of a search parameter space and possible signal distributions over it.

#### A. Simulated data and trigger generation

The space of “low-mass” binaries with non-spinning components of mass between 1–24  $M_\odot$  and total mass up to 25  $M_\odot$  was the basis of a search in recent LIGO-Virgo data [32]. Using standard LAL codes we generated template banks at 2PN order

for this search space, and obtained noise background triggers by matched filtering approximately 18 days of Gaussian coloured noise with a sensitivity curve (PSD) representative of the ‘‘Early Advanced LIGO’’ operational sensitivity shown in [38], for the LIGO Hanford (LHO) and LIGO Livingston (LLO) detectors. We used the stationary-phase approximation (SPA) filters of [11], corresponding to the TaylorF2 approximant, with a lower frequency cutoff of 40 Hz, and approximately 11300 templates were placed for each detector. The 2PN order was chosen for simplicity given that the appropriate metric [16] and template placement algorithm are well-established; metric-based template banks for (hopefully) more accurate waveform models are just becoming available [6, 8].

For simplicity we simulated an astrophysical distribution of signal events with the property that the product  $R(m_1, m_2)\mathcal{M}^{5/2}$  is constant over the search space, implying a constant rate density of signals at fixed SNR per  $dm_1 dm_2$  element. We note that this distribution implies an astrophysical rate which is a *negative* power of mass; this is consistent with the simple expectation that the distribution of compact objects may mirror a stellar initial mass function up to some maximum mass. The signal distribution was implemented by generating a set of simulated signals with the TaylorT2 approximant [36, 39] at 2PN order with random parameters distributed uniformly over the component masses and over angular parameters, then rescaling the distances to achieve a  $1/\bar{\rho}^4$  distribution of expected SNRs, above a minimum value chosen below the search threshold  $\rho_{\text{th}} = 5.5$ , and up to a maximum  $\bar{\rho}$  value well above any noise event<sup>10</sup>.

To reduce the number of random noise triggers in our signal event sample we also imposed the parameter consistency test of [30] between triggers in the two LIGO detectors. Despite the application of clustering, there may occasionally be more than one coincident trigger recovered near the coalescence time of a simulated signal; in these cases we took the trigger with the highest  $\rho$  within  $\pm 0.025$  s of that time to represent the signal.<sup>11</sup> We thus obtained approximately 8000 simulated signal events.

Since our signal recovery involved choosing the loudest trigger by SNR over an 0.05 s window around the simulated signal time, to treat noise triggers consistently we also clustered them over 0.05 s time windows. This extra clustering step had a negligible effect on triggers above a SNR of approximately 6.0 and we found our results were not affected by it except at very high false alarm rates (FARs). We obtained approximately  $5.2 \times 10^6$  clustered noise triggers per detector.

## B. Evaluation of the optimized statistic

Then, considering only triggers from a single detector’s data, we evaluated the optimized statistic given by the RHS of Eq. (26) for both background and simulated signal triggers, using  $m_1, m_2$  coordinates for the signal rate and metric. We modelled the noise trigger distribution  $p(\rho|N)$  above threshold by a falling exponential  $\propto e^{-\alpha\rho}$ , finding  $\alpha \simeq \ln(250)$ . The metric components are pre-computed to 2PN order in  $(t_c, \tau_0, \tau_3)$  coordinates; the density in  $(t_c, m_1, m_2)$  space is then given by  $\sqrt{\det \gamma(t_c, m_1, m_2)} \equiv |\det J| \sqrt{\det \gamma(t_c, \tau_0, \tau_3)}$ , where  $J$  is the Jacobian with

$$\det J = \det \begin{bmatrix} \frac{\partial \tau_0}{\partial m_1} & \frac{\partial \tau_3}{\partial m_1} \\ \frac{\partial \tau_0}{\partial m_2} & \frac{\partial \tau_3}{\partial m_2} \end{bmatrix} = \frac{A_0 A_3}{m_1^3 m_2^3} (m_1 + m_2)^{2/3} (m_1 - m_2), \quad (27)$$

where we omitted trivial elements involving  $t_c$  (unity on the diagonal and zero elsewhere).

*Regularization and edge effects* The resulting density  $\sqrt{\det \gamma(t_c, m_1, m_2)}$  vanishes on the  $m_1 = m_2$  line, which would naïvely assign an *infinite* statistic value to any triggers in equal-mass templates. This singularity presages the breakdown of our approximation of a constant prior over a peak of likelihood near this line. The width of such a peak, at a given SNR, is approximately constant in  $(\tau_0, \tau_3)$  coordinates; in these coordinates the prior density is multiplied by the inverse of the Jacobian  $J$ , which blows up at  $m_1 = m_2$ .<sup>12</sup> Thus, for SNR maxima close to this line the appropriate optimal statistic is not given by Eq. (15), rather it would require an explicit integration with a varying prior, for which the result would not diverge.

In addition to this singularity we will in general meet artefacts due to the discreteness of the template bank: the search parameter space consists of cells of finite area centred on each template, but we only evaluate the metric at the template itself. For a given template, the  $\Lambda_{\text{opt}}$  value assigned to triggers should ideally arise from an integral over the area covered by the template. In practice, artefacts due to finite cell size appear to be negligible except again near the  $m_1 = m_2$  line, where the re-weighting factor varies strongly from one template to the next, and for templates close to the large- $\mathcal{M}$  edge of the bank, mostly those with  $m_1, m_2 > 7.5 M_\odot$  (see Fig. 1). Such templates may be ranked anomalously highly considering that the areas they cover extend beyond the limits of the bank (at  $\eta = 0.25$  and  $m_1 + m_2 = 25M_\odot$ ) where there are no simulated signals.

Naïvely applying the inverse metric formula (26) without regard to these caveats severely hurts the efficiency of our statistic, even if we ‘‘remove’’ the singularity by making sure no template has exactly equal masses; we find that statistic values in noise are

<sup>10</sup> More precisely, the network SNR for the LHO and LLO detectors  $(\bar{\rho}_H^2 + \bar{\rho}_L^2)^{1/2}$  was rescaled to the target  $1/\bar{\rho}^4$  distribution up to a maximum of 20; virtually all the loudest signals in this distribution are well above any (Gaussian) noise event in either detector.

<sup>11</sup> This choice may be sub-optimal when using a ranking statistic other than  $\rho$ ; we should instead use the trigger with the best statistic value associated with the simulated signal. Doing so would only strengthen any gain in efficiency due to an improved statistic.

<sup>12</sup> Equivalently, in  $(m_1, m_2)$  coordinates the prior is constant but the *area* of integration in parameter space covered by a peak of likelihood, at least in the linear signal approximation, blows up at equal mass.




FIG. 2. False alarm rate (FAR) values assigned to signals with given matched filter SNR  $\rho$ . We show FARs assigned using different ranking statistics, either  $\rho$  (solid black line) or the likelihood ratio  $\Lambda_{\text{opt}}$  for a prior distribution of signals uniform over  $(m_1, m_2)$  at constant SNR (dashed/dotted lines); for  $\Lambda_{\text{opt}}$  we show the FARs assigned for various possible signal component mass values, as in the legend.

dominated by a few triggers from the region  $m_1, m_2 > 7.5 M_{\odot}$ , causing the number of signals recovered at low false alarm rates to decline drastically. Clearly such templates, although they may be relatively likely to see signals, are not correctly described by our first approximation.

We choose to regularize the statistic values assigned to such templates by adding a small constant in quadrature to  $\sqrt{\det \gamma(t_c, m_1, m_2)}$ , such that the distribution of  $\Lambda_{\text{opt}}$  values in noise is not dominated by outliers from templates with small metric densities. Having done so, we no longer see a decline in search efficiency at low false alarm rates. We tested different choices of regularization constant and found that efficiency was optimised within a range of values for which  $< 1\%$  of noise triggers were significantly affected by regularization, and was insensitive to variation within this range. As remarked earlier, a more theoretically correct approach would be to model the variation of the prior over peaks of likelihood, including edge effects, in appropriate coordinates; however this would be significantly more complicated to implement and we see no evidence that doing so would lead to significant improvements in performance.

### C. Improvement in search efficiency

For each simulated signal, we calculate the false alarm rate (FAR) as the number out of the 5.2 million noise events assigned a ranking statistic (either matched filter  $\rho$  or our optimized statistic  $\Lambda_{\text{opt}}$ ) equal to or higher than that of the signal trigger, divided by the search time.<sup>13</sup>

We illustrate the relation between signal SNR and the resulting FAR value in Figure 2, where for the likelihood ratio statistic we have chosen some representative points in the template bank to illustrate the effect of mass-dependent re-weighting. There is a non-trivial difference between the matched filter values for different candidate masses that result in the same FAR; thus, the search has a higher sensitive volume (at fixed FAR) in regions of parameter space where the rate of signals relative to noise events is higher, and vice versa. Our mass-dependent statistic leads to an overall improvement in search efficiency since the gain in the expected number of signals from regions where the sensitivity is increased (relative to ranking by SNR) outweighs the loss from regions where it decreases.

The net effect of the re-weighting is thus to increase the number of simulated signals seen with a given FAR value or lower. We show this improvement in Figure 3 where we see that the search efficiency, for the given parameter space and signal distribution, is improved by some tens of percent. Note that the evaluation of such ROCs is subject to fluctuations due to small number statistics at the extreme low FAR end. This result can also be interpreted in terms of a reduction in false alarm rate: by using the optimized statistic, a given number of signals is found at or below a FAR which is approximately a factor 10 lower than when using  $\rho$  as ranking statistic.

Interestingly, if instead of the 2PN metric density we use the density of a 2-d metric with constant components over  $\tau_0, \tau_3$  space to evaluate  $\Lambda_{\text{opt}}$ , we obtain almost the same ROC curve with only a small loss of efficiency relative to the theoretically correct

<sup>13</sup> Note that the resolution of the FAR value is limited by the search time; much smaller FARs can be assigned by multi-detector searches.




FIG. 3. Number of simulated signals recovered at a given false alarm rate (FAR) using either the matched filter SNR, or the optimized statistic  $\Lambda_{\text{opt}}$  constructed using the 3-d metric  $\gamma(t_c, m_1, m_2)$  for 2PN inspiral filters, as ranking statistic.

2PN metric at low FAR. The great majority of the improvement over SNR ranking is thus due to the strong mass-dependence of the Jacobian of Eq. (27).

We performed an independent check for the optimal classification of noise and simulated signal triggers by constructing a random forest classifier [40, 41] with stratified cross-validation, where both sets of triggers are divided into 5 equally-sized “folds” and the probability of signal *vs.* noise is evaluated for triggers in each fold using a forest trained on the remaining four, using the matched filter SNR and (simple functions of) the binary component masses as classification features. Using the forest predicted probability as ranking statistic, we found the efficiency as a function of FAR was within 2-3% of that obtained by our  $\Lambda_{\text{opt}}$  statistic for FAR values above a few  $\times 10^{-6}$  Hz; at lower FARs the random forest had somewhat worse efficiency, possibly due to the difficulty of training in regions of feature space with very few noise triggers. This result gives us more confidence that the statistic we have implemented is close to optimal, at least within the statistical limitations of our example. Note that, even after imposing a cut on SNR to reduce the total trigger count, training a classifier with 200 trees required order(1/2 hour) computing time compared to a few minutes for evaluation of  $\Lambda_{\text{opt}}$  using pre-computed metric components for all 5.2 million noise triggers.

#### D. Effect of varying signal distributions

A natural question is how far the benefits produced by using an optimized statistic rather than the matched filter SNR are sensitive to a discrepancy between the true distribution of signals in nature and the distribution assumed in constructing the statistics. Note too that the goal of this paper is to describe a *method* of optimizing the ranking statistic of a search for *any* given prior signal distribution; we make no claims that the prior uniform over  $(m_1, m_2)$  at fixed SNR, which we chose for purposes of illustration, is in fact optimal for the true astrophysical distribution of binaries.

We can address the issue of discrepancies in the assumed *vs.* real signal distributions either by changing the prior  $R(\vec{M})D_{\text{hor}}(\vec{M})^3$  used in evaluating the  $\Lambda_{\text{opt}}$  statistic, but keeping the same set of simulated signals to evaluate efficiency; or by keeping the dependence of the ranking statistic on  $\vec{M}$  unchanged, but changing the weighting of simulated signals (or using a different set of simulations entirely) to correspond to a different true distribution. Many different possible distributions of binary coalescence parameters have been proposed (see [42] for a recent study) and we will only consider a few simple scenarios here.

Figure 4 shows the results of our comparisons. The left plot shows that, for a population of signals uniform over  $(m_1, m_2)$ , the statistic obtained by optimizing for a distribution uniform over  $(\log m_1, \log m_2)$  shows almost the same gain in efficiency as the “correctly” optimized statistic for the true signal population.<sup>14</sup> In the right plot we evaluate the ROC by totalling the simulated signals recovered with a FAR below threshold, with a parameter-dependent weight proportional to  $1/(m_1 m_2)$ ; the resulting figure of merit is appropriate to a true signal population uniform on  $(\log m_1, \log m_2)$ . We see in this case, as expected,

<sup>14</sup> The optimization in either case is not perfect, due to the lower PN order used for the metric evaluation compared to the filters, and other approximations made in deriving the statistic.

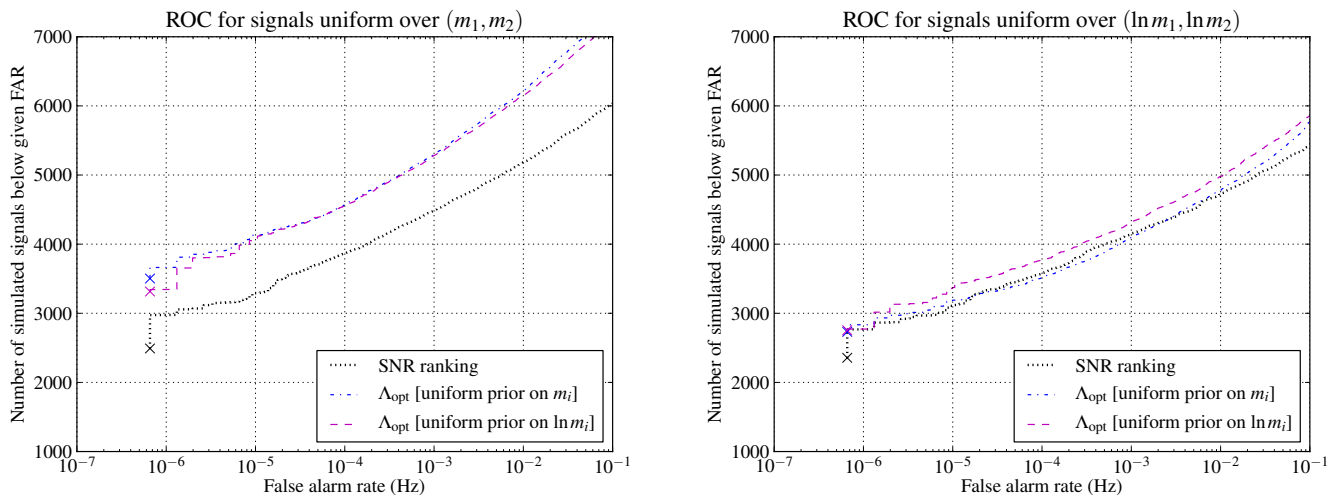


FIG. 4. *Left*—Comparison of matched filter SNR and likelihood statistics constructed using different assumed prior signal distributions (either uniform in  $(m_1, m_2)$  or uniform in  $(\log m_1, \log m_2)$ ); ROCs are calculated with the injected signal distribution which is uniform over  $m_1, m_2$ . *Right*—Comparison of statistics, as in the left panel: here, ROCs are calculated by re-weighting the injected signal distribution to simulate a population uniform over  $(\log m_1, \log m_2)$ .

the statistic optimized using the “correct” prior, i.e. uniform over  $(\log m_1, \log m_2)$ , has the best efficiency; the statistic optimized for a signal distribution uniform over  $(m_1, m_2)$  performs somewhat worse, but not significantly worse than the matched filter SNR. The differences in efficiency are not large for this case; note that the signal distribution uniform on  $(\log m_1, \log m_2)$  is intermediate between the metric density (i.e. the prior distribution implied by ranking on SNR) and the uniform-over-component-mass distribution, in terms of how strongly lower-mass signals are favoured.

More extreme “worse case” scenarios can be conceived where the true signal population deviates further from the prior used to evaluate the  $\Lambda_{\text{opt}}$  statistic, and we would expect the efficiency of a search using such an “unrealistic” prior to suffer accordingly. The limiting case would be for the entire population of real signals to be concentrated at the point where the assumed prior is smallest. There would, nevertheless, be a nonzero lower bound on the efficiency of the search, provided that the prior did not vanish at that point.

An example of a distribution which differs very strongly from an assumed prior uniform over component masses is one where *only binary neutron star* (BNS) coalescences occur, i.e. black hole binaries are entirely absent.<sup>15</sup> As shown in Figure 2, the SNR required for a BNS of masses (1.35, 1.3) to be assigned a given false alarm rate below  $\sim 10^{-5}$ , using our statistic optimized for a uniform mass distribution, is greater than that required using SNR as ranking statistic, but only by approximately 0.5. Thus in this case, using a statistic optimised for a uniform mass distribution (including black hole binaries) will perform worse than the implied prior from the template metric, detecting a factor of approximately  $(8/8.5)^3 \approx 0.83$  fewer signals at low FAR. Detailed discussions of the “best” choice of search prior given widely differing possible true astrophysical distributions lie beyond the scope of this paper, but we make a few further remarks in the concluding Section.

## V. DISCUSSION

We have shown that the optimal statistic for detection of transient gravitational wave signals of known form over a broad intrinsic parameter space is *not* the matched filter SNR, even in Gaussian noise, so long as the distribution of signals over the parameter space is not identical to the metric density associated with the matched filter. Since the signal distribution is determined by astrophysical processes, it is unlikely to be close to this default distribution; thus, an optimized Bayes factor detection statistic can in general outperform the matched filter SNR. We showed how to approximately evaluate this statistic in a search for coalescing binaries with non-spinning components, given only the information contained in “triggers” output by the standard LIGO inspiral search code. Extending the method to searches for binaries with non-precessing component spins is straightforward, given the recent development of metric-based methods [6] for covering the higher-dimensional space of intrinsic parameters required in this case.

<sup>15</sup> This example was suggested by the referee.

*Waveform uncertainty and parameter bias* In the proof-of-principle analysis performed in this paper we have chosen to simulate signals using the TaylorT2 approximant and perform the matched filter search using a stationary-phase approximation akin to TaylorF2 [39], both at second post-Newtonian order. This choice ensures that the parameter bias on the maximum likelihood mass template will be relatively small [35, 36]. No choice of model template for a search will correspond precisely to the true gravitational waveform, so in general the distribution of recovered signal triggers in intrinsic parameter space will not correspond to the true distribution of sources. In addition to reductions in SNR caused by an imperfect fitting factor between the true waveform and the filter, intrinsic parameter bias will reduce the efficiency of the search if it leads to an inaccurate assignment of ranking statistic values and if the statistic is optimized for the true signal distribution rather than the recovered one. However, the loss of efficiency caused by parameter bias is probably less important than the sub-optimality caused by our lack of knowledge of the true astrophysical distribution (see below).

For the non-spinning template bank used in our example, bias in parameter recovery might not greatly affect the efficiency of our statistic:<sup>16</sup> the errors on recovered chirp mass  $\mathcal{M}$  are expected to be very small, and it may be verified that the metric density (in component mass coordinates) varies weakly along lines of constant  $\mathcal{M}$  except for a small region near the equal-mass line.

*Alternative marginalization methods* An alternative path to evaluating the Bayes factor (marginalized likelihood ratio) would be to use the template bank as a set of sampling points to explicitly evaluate the integral over intrinsic parameters  $\vec{M}$  for short time intervals and restricted regions of parameter space. Thus, instead of taking the maximum of SNR in a trigger cluster to represent the entire peak of likelihood, one would add up the likelihood contributions over templates in a cluster and record the total, multiplying by the prior density (in appropriate coordinates) to obtain the analogue of Eq. (26). Whether this method is preferable to approximating the integrand by a multivariate Gaussian will depend on technical issues such as the density of the bank, the accuracy to which it fulfils the minimal match criterion and the accuracy to which the metric is known. For instance explicit integration over templates may be more suitable for a search with phenomenological inspiral-merger-ringdown filters where a metric cannot readily be calculated.

*Multi-detector search* Our method leads to an optimized statistic when considering parameters affecting the signal seen in a single interferometric detector. The re-weighting procedure over the “intrinsic” parameters (masses and spins) remains essentially the same when considering more than one (spatially separated) detector. However, additional information is available from a multi-detector search, namely the differences in coalescence time and phase, and the relative signal amplitudes between detectors. The signal population will have a (possibly strongly) different distribution from that of noise events over these parameters; thus, knowledge of these distributions will help in distinguishing between them, and should further improve search efficiency relative to the standard multi-detector statistic (the quadrature sum of single-detector matched filter values [12]).

Exploiting this information is a more complex task which we leave to future work. Machine learning classification appears the most practical strategy to deal with this multi-dimensional problem; in [43], a “naïve Bayes” framework was used to derive a ranking statistic using measured distributions over coincident event parameters for a modelled gravitational-wave burst search, although correlations between these parameters were not modelled.

*Non-Gaussian noise* The presence of loud, unpredictable non-Gaussian artefacts (“glitches”) in the detector output [44] is a major challenge for attempts to calculate an optimal detection statistic from first principles. Various ad-hoc signal consistency tests have been implemented [45, 46] and shown to be effective in real noise at separating simulated signals from glitches, however they require tuning “by hand” to the data, and are not immediately connected to the Bayesian framework described here. A theoretical approach to searching data containing outlying non-Gaussian samples was outlined in [47], while the use of an explicit glitch model in evaluating Bayes factors was explored in [48]; however so far an implementation suitable to search extended periods of data for coalescing binary signals is not available.

One approach using the existing trigger values calculated in the inspiral search algorithm, including the signal consistency tests, would be to simply treat these as additional parameters and attempt to estimate (fit) the relative signal and noise distributions empirically, along the lines of Section III C.<sup>17</sup> The result would correspond to the Bayes factor calculated over some subset of the data corresponding to the quantities available. However, this might miss a considerable amount of information that could have helped exclude glitches; thus, the standard chi-squared test [45] does not distinguish between high- and low-frequency transients, although they may have very different distributions in any given stretch of data.

No single method can be optimal for data containing arbitrary non-Gaussian artefacts, since the space of possible non-Gaussian behaviours is infinite; at most we can look for methods that approach optimality for populations of glitches resembling those seen in real data. Note that any such method must incorporate our findings on the density of Gaussian noise triggers, since it must reduce to the optimal statistic described here in the Gaussian noise limit (for a single detector).

*Choice of prior and astrophysical uncertainties* In the Introduction we mentioned the very large current uncertainties in the astrophysical distributions of coalescing binary systems. The question is then, what interpretation can we give to the prior distribution over intrinsic parameters that enters our proposed optimal statistic. Two cases can be distinguished: pre-detection,

<sup>16</sup> We in fact started our case study using TaylorT4 simulated signal waveforms, which have a significant bias relative to the TaylorF2 templates; despite this we obtained a comparable gain in search efficiency.

<sup>17</sup> Ref. [20] describes an approach similar in spirit, but without explicitly modelling the trigger distributions.

where no significant additional information from GW observations is available;<sup>18</sup> and post-detection, where the prior for future searches will be informed by the recovered parameter values of some number of binary signals.

In the pre-detection case one might choose the optimization prior to correspond to an astrophysical model of compact binary formation and evolution (for instance [42]), after having marginalized over unknown nuisance parameters of the model. However, this approach has potential pitfalls, since one might mistakenly give a negligibly small weight to systems that the model (wrongly) predicts to be extremely rare, and thus produce a “blind spot” in the search; also, the complexity of such models is likely to make the search results difficult to interpret. Alternatively, one might want to verify the performance of an analysis pipeline before deployment by checking its performance in recovering simulated signals, *i.e.* a Monte Carlo or mock data challenge, and demanding that the analysis be optimized for a well-defined signal population covering the intended parameter space [49]. As pointed out by [18], for every analysis choice there may be *some* signal distribution for which it is optimal, but this distribution may be either unphysical, implausible, or fail to correspond to the search requirements; use of an optimized statistic can avoid such issues.

Post-detection, it will be possible to rule out some signal distributions, thus the range of possible search priors will decrease. How to proceed in detail will depend on the scientific objectives: for instance one might aim to maximize the number of detected signals, or to obtain the best accuracy on measurements of astrophysical coalescence rates as a function of source parameters, or to obtain the largest Bayes factor in distinguishing between different models of compact binary formation and evolution. We will leave such more complex optimization problems for future investigations.

### ACKNOWLEDGEMENTS

We would like to thank Drew Keppel, Reinhard Prix, Andrew Lundgren, Ian Harry and Walter Del Pozzo for enlightening discussions; Will Farr for carefully reading an earlier draft of this paper; the LIGO Scientific Collaboration for the generation and storage of the standard Gaussian data set used in this study; and the Albert-Einstein-Institut (Hannover) for the use of the Atlas computing cluster. JV was supported by the research programme of the Foundation for Fundamental Research on Matter (FOM), which is partially supported by the Netherlands Organisation for Scientific Research (NWO); TD acknowledges support from the Max-Planck-Gesellschaft.

### Appendix A: Marginalization over phase

Here we will derive the phase-marginalized likelihood function and expand it around a maximum. The full likelihood including phase is written in terms of the two phases of the signal  $\hat{h}_0$  and  $\hat{h}_{\pi/2}$ :

$$L = \exp\left(-\frac{\mathcal{A}^2}{2} \langle \hat{h}_0 | \hat{h}_0 \rangle + \mathcal{A} \langle \hat{h}_0 | d \rangle \cos \phi_c + \mathcal{A} \langle \hat{h}_{\pi/2} | d \rangle \sin \phi_c\right). \quad (\text{A1})$$

We marginalize over  $\phi_c$  using a uniform prior  $p(\phi_c|S) = (2\pi)^{-1}$  via the integral

$$L' = \exp\left(-\frac{\mathcal{A}^2}{2} \langle \hat{h} | \hat{h} \rangle\right) \int_0^{2\pi} \frac{d\phi_c}{2\pi} \exp\left(\mathcal{A} [\langle \hat{h}_0 | d \rangle \cos \phi_c + \langle \hat{h}_{\pi/2} | d \rangle \sin \phi_c]\right), \quad (\text{A2})$$

which we perform by substituting  $z = e^{i\phi_c}$  and  $a = \mathcal{A} \langle \hat{h}_0 | d \rangle + i\mathcal{A} \langle \hat{h}_{\pi/2} | d \rangle$  to arrive at a standard result in terms of the modified Bessel function  $I_0$ :

$$L' = \frac{1}{2\pi} \exp\left(-\frac{\mathcal{A}^2}{2} \langle \hat{h} | \hat{h} \rangle\right) \oint dz \exp\left(\frac{1}{2}(za + z^{-1}a^*)\right) \quad (\text{A3})$$

$$= \exp\left[-\frac{\mathcal{A}^2}{2} \langle \hat{h} | \hat{h} \rangle + \log I_0\left(\mathcal{A} \sqrt{\langle \hat{h}_0 | d \rangle^2 + \langle \hat{h}_{\pi/2} | d \rangle^2}\right)\right]. \quad (\text{A4})$$

<sup>18</sup> Upper limits from GW non-observations could *somewhat* restrict the range of possible priors.



## 1. Maximum likelihood location

In general, the parameters which maximize  $L$  can be different from those that maximize  $L'$ , so we need to recalculate them to check that our use of the matched filter maximum is still valid. At the maximum we have

$$0 = \partial_\alpha \left[ -\frac{\mathcal{A}^2}{2} \langle \hat{h} | \hat{h} \rangle \right] + \frac{I_1(\mathcal{A}\rho(\hat{h}))}{I_0(\mathcal{A}\rho(\hat{h}))} \partial_\alpha [\mathcal{A}\rho(\hat{h})], \quad (\text{A5})$$

where  $\rho(\hat{h}) \equiv \sqrt{\langle \hat{h}_0 | d \rangle^2 + \langle \hat{h}_{\pi/2} | d \rangle^2}$  and  $\alpha$  indexes all parameters except  $\phi_c$ , including the amplitude. First we consider the amplitude  $\mathcal{A}$ , where we find the maximum at  $\mathcal{A}'_{\text{ML}}$  given by the solution of

$$\mathcal{A}'_{\text{ML}} = \frac{I_1(\mathcal{A}'_{\text{ML}}\rho(\hat{h}^{\text{ML}}))}{I_0(\mathcal{A}'_{\text{ML}}\rho(\hat{h}^{\text{ML}}))} \rho(\hat{h}^{\text{ML}}), \quad (\text{A6})$$

such that

$$\lim_{\mathcal{A}'_{\text{ML}} \rightarrow \infty} \mathcal{A}'_{\text{ML}} = \rho(\hat{h}^{\text{ML}}) = \mathcal{A}_{\text{ML}}. \quad (\text{A7})$$

As the ratio of modified Bessel functions tends to 1 for above-threshold detection statistic values (e.g.  $I_1(36)/I_0(36) \approx 0.986$  for SNR=6), in this limit we recover the same maximum likelihood point as for the unmarginalized  $L$  function. However, in the low SNR limit the recovered maximum  $\mathcal{A}'_{\text{ML}} < \mathcal{A}_{\text{ML}}$ .

For the intrinsic parameters labelled by  $p$  we use Eq. (A5) at the maximum likelihood point to find

$$\partial_p \log L'_{\text{ML}} = -\mathcal{A}_{\text{ML}}^2 \langle \partial_p \hat{h}_0 | \hat{h}_0 \rangle + \frac{I_1(\mathcal{A}_{\text{ML}}^2)}{I_0(\mathcal{A}_{\text{ML}}^2)} \left[ \langle \hat{h}_0 | d \rangle \langle \partial_p \hat{h}_0 | d \rangle + \langle \hat{h}_{\pi/2} | d \rangle \langle \partial_p \hat{h}_{\pi/2} | d \rangle \right]. \quad (\text{A8})$$

Using the normalisation condition which implies  $\langle \partial_p \hat{h} | \hat{h} \rangle = 0$  for all parameter values, we also find

$$0 = \partial_p \partial_q \langle \hat{h} | \hat{h} \rangle = \langle \partial_p \partial_q \hat{h} | \hat{h} \rangle + \langle \partial_p \hat{h} | \partial_q \hat{h} \rangle, \quad (\text{A9})$$

$$\langle \hat{h}_0 | d \rangle \langle \partial_p \hat{h}_0 | d \rangle_{\text{ML}} = -\langle \hat{h}_{\pi/2} | d \rangle \langle \partial_p \hat{h}_{\pi/2} | d \rangle_{\text{ML}} \implies \partial_p \rho(\hat{h})_{\text{ML}} = 0. \quad (\text{A10})$$

## 2. Expansion around maximum

In order to expand  $\log L'$  we also need the second derivatives with respect to the parameter values. For the amplitude term  $\partial_{\mathcal{A}} \partial_{\mathcal{A}}$  we have

$$\partial_{\mathcal{A}} \partial_{\mathcal{A}} \log L' \approx -1 + \mathcal{A}_{\text{ML}}^2 \left( 1 + \frac{I_2(\mathcal{A}_{\text{ML}}^2)}{I_0(\mathcal{A}_{\text{ML}}^2)} - 2 \frac{I_1(\mathcal{A}_{\text{ML}}^2)^2}{I_0(\mathcal{A}_{\text{ML}}^2)^2} \right), \quad (\text{A11})$$

where the second term tends to zero at high SNR, so  $\partial_{\mathcal{A}} \partial_{\mathcal{A}} \log L' \approx -1$ . For the intrinsic parameters we obtain

$$\partial_p \partial_q \log L'_{\text{ML}} = \frac{I_1(\mathcal{A}_{\text{ML}}^2)}{I_0(\mathcal{A}_{\text{ML}}^2)} \cdot \left[ \langle \hat{h}_0 | d \rangle \langle \partial_p \partial_q \hat{h}_0 | d \rangle + \langle \hat{h}_{\pi/2} | d \rangle \langle \partial_p \partial_q \hat{h}_{\pi/2} | d \rangle + \langle \partial_p \hat{h}_0 | d \rangle \langle \partial_q \hat{h}_0 | d \rangle + \langle \partial_p \hat{h}_{\pi/2} | d \rangle \langle \partial_q \hat{h}_{\pi/2} | d \rangle \right]_{\text{ML}}, \quad (\text{A12})$$

where derivatives of the Bessel functions vanish at the ML point via (A10), which also implies  $\partial_{\mathcal{A}} \partial_p \log L' = 0$  at this point. To proceed we write

$$d = \mathcal{A}_{\text{ML}} \hat{h}^{\text{ML}} + n = \mathcal{A}_{\text{ML}} (\cos \phi_c^{\text{ML}} \hat{h}_0 + \sin \phi_c^{\text{ML}} \hat{h}_{\pi/2}) + n \quad (\text{A13})$$

where  $\tan \phi_c^{\text{ML}} = \langle \hat{h}_{\pi/2}^{\text{ML}} | d \rangle / \langle \hat{h}_0^{\text{ML}} | d \rangle$  and  $n$  is a noise residual having vanishing overlap with  $\hat{h}^{\text{ML}}$ , which is formally of sub-leading order in  $\mathcal{A}_{\text{ML}}$  and can be neglected for a sufficiently high SNR [26].<sup>19</sup> We then find, after some rearrangement and applying

<sup>19</sup> Note that this may not hold in the presence of a signal which deviates significantly from the model  $\hat{h}(\vec{\theta})$ .

(A10),

$$\begin{aligned} \partial_p \partial_q \log L'_{\text{ML}} &\simeq \frac{I_1(\mathcal{A}_{\text{ML}}^2)}{I_0(\mathcal{A}_{\text{ML}}^2)} \left[ -\mathcal{A}_{\text{ML}}^2 \langle \partial_p \hat{h}_0 | \partial_q \hat{h}_0 \rangle + \frac{\langle \partial_p \hat{h}_0 | d \rangle \langle \partial_q \hat{h}_0 | d \rangle}{\sin^2 \phi_c^{\text{ML}}} \right]_{\text{ML}} \\ &\simeq -\mathcal{A}_{\text{ML}}^2 \frac{I_1(\mathcal{A}_{\text{ML}}^2)}{I_0(\mathcal{A}_{\text{ML}}^2)} \gamma_{pq}(\hat{h}^{\text{ML}}), \end{aligned} \quad (\text{A14})$$

where  $\gamma_{pq} = \langle \partial_p \hat{h}_0 | \partial_q \hat{h}_0 \rangle - \langle \partial_p \hat{h}_0 | \hat{h}_{\pi/2} \rangle \langle \partial_q \hat{h}_0 | \hat{h}_{\pi/2} \rangle$  is the *phase-projected metric* describing overlaps maximized over phase difference. We then find the expansion of the likelihood

$$\log L'(\mathcal{A}_{\text{ML}} + \Delta\mathcal{A}, \vec{\theta}^{\text{ML}} + \Delta\vec{\theta}) \simeq -\frac{\mathcal{A}_{\text{ML}}^2}{2} - \frac{\Delta\mathcal{A}^2}{2} + \log I_0(\mathcal{A}_{\text{ML}}^2) + \frac{1}{2} \Delta\theta_p \Delta\theta_q (\partial_p \partial_q \log L')_{\text{ML}}. \quad (\text{A15})$$

With the assumption that the priors are slowly varying around the maximum, we obtain, for an n-dimensional parameter vector  $\vec{\theta}$ ,

$$\int d\mathcal{A} d^n \vec{\theta} p(\mathcal{A}, \vec{\theta} | S) L' \simeq p(\mathcal{A}, \vec{\theta} | S)_{\mathcal{A}_{\text{ML}}, \vec{\theta}^{\text{ML}}} L'_{\text{ML}} \int d(\Delta\mathcal{A}) e^{-\Delta\mathcal{A}^2/2} \int d^n \vec{\theta} \exp \left[ -\frac{1}{2} \mathcal{A}_{\text{ML}}^2 \gamma'_{pq} \Delta\theta_p \Delta\theta_q \right] \quad (\text{A16})$$

$$= p(\mathcal{A}_{\text{ML}}, \vec{\theta}^{\text{ML}} | S) \exp \left( -\frac{\mathcal{A}_{\text{ML}}^2}{2} \right) I_0(\mathcal{A}_{\text{ML}}^2) \mathcal{A}_{\text{ML}}^{-n} \sqrt{\frac{(2\pi)^{n+1}}{\det \gamma'(\vec{\theta}^{\text{ML}})}}, \quad (\text{A17})$$

where  $\gamma'$  is equal to  $\gamma$  rescaled by the ratio of Bessel functions appearing in Eq. (A14); for moderate or large SNRs this ratio is very close to unity, hence we may approximate  $\gamma' \simeq \gamma$ .

- 
- [1] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Class. Quant. Grav.* **27**, 173001 (2010), arXiv:1003.2480.
- [2] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash, *Class. Quant. Grav.* **23**, 5477 (2006), arXiv:gr-qc/0604037.
- [3] T. Cokelaer, *Phys. Rev. D* **76**, 102004 (2007), arXiv:0706.4437.
- [4] D. Keppel, A. P. Lundgren, B. J. Owen, and H. Zhu, *Phys. Rev.* **D88**, 063002 (2013), arXiv:1305.5381.
- [5] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, *Phys. Rev.* **D86**, 064020 (2012), arXiv:1203.6060.
- [6] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, *Phys. Rev.* **D86**, 084017 (2012), arXiv:1207.6406.
- [7] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. Weinstein, (2012), arXiv:1210.6666.
- [8] I. Harry, A. Nitz, D. A. Brown, A. Lundgren, E. Ochsner, *et al.*, (2013), arXiv:1307.3562.
- [9] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, *et al.*, *Phys. Rev.* **D89**, 024003 (2014), arXiv:1310.5633 [gr-qc].
- [10] D. A. Brown, *Search for gravitational radiation from black hole MACHOs in the Galactic halo*, Ph.D. thesis, University of Wisconsin-Milwaukee (2004), arXiv:0705.1514 [gr-qc].
- [11] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [12] S. Babak *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [13] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev.* **D79**, 122001 (2009), arXiv:0901.0302.
- [14] R. Balasubramanian, B. Sathyaprakash, and S. Dhurandhar, *Phys. Rev.* **D53**, 3033 (1996), arXiv:gr-qc/9508011.
- [15] B. J. Owen, *Phys. Rev.* **D53**, 6749 (1996), arXiv:gr-qc/9511032.
- [16] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [17] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
- [18] A. C. Searle, P. J. Sutton, and M. Tinto, *Class. Quant. Grav.* **26**, 155017 (2009), arXiv:0809.2809.
- [19] R. Prix and B. Krishnan, *Class. Quant. Grav.* **26**, 204013 (2009), arXiv:0907.2569.
- [20] R. Biswas, P. R. Brady, J. Burguet-Castell, K. Cannon, J. Clayton, *et al.*, *Phys. Rev.* **D85**, 122008 (2012), arXiv:1201.2959.
- [21] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev.* **D88**, 024025 (2013), arXiv:1209.0718.
- [22] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992), arXiv:gr-qc/9209010.
- [23] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [24] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010).
- [25] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **80**, 063007 (2009).
- [26] E. Poisson and C. M. Will, *Phys. Rev.* **D52**, 848 (1995), arXiv:gr-qc/9502040.
- [27] M. Vallisneri, *Phys. Rev.* **D77**, 042001 (2008), arXiv:gr-qc/0703086.
- [28] N. J. Cornish and T. B. Littenberg, *Phys. Rev. D* **76**, 083006 (2007).
- [29] C. Röver, *Journal of Physics Conference Series* **228**, 012008 (2010), arXiv:0911.5051.
- [30] C. Robinson, B. Sathyaprakash, and A. S. Sengupta, *Phys. Rev.* **D78**, 062002 (2008), arXiv:0804.4816.
- [31] B. F. Schutz, *Class. Quant. Grav.* **28**, 125023 (2011), arXiv:1102.5421.

- [32] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys. Rev.* **D85**, 082002 (2012), arXiv:1111.7314.
- [33] C. Robinson, B. Sathyaprakash, and A. Sengupta, *Taking parameter correlations into account in the binary inspiral pipeline*, LIGO-DCC G070460-00 (LIGO, 2007).
- [34] S. Fairhurst and P. Brady, *Class.Quant.Grav.* **25**, 105002 (2008), arXiv:0707.2410.
- [35] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [36] A. H. Nitz, A. Lundgren, D. A. Brown, E. Ochsner, D. Keppel, and I. W. Harry, *ArXiv e-prints* (2013), arXiv:1307.1757.
- [37] <http://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [38] J. Aasi *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2013), arXiv:1304.0670.
- [39] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen, *Phys. Rev. D* **59**, 124016 (1999).
- [40] [http://www.stat.berkeley.edu/~breiman/RandomForests/cc\\_home.htm](http://www.stat.berkeley.edu/~breiman/RandomForests/cc_home.htm).
- [41] F. Pedregosa *et al.*, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [42] M. Dominik *et al.*, *The Astrophysical Journal* **759**, 52 (2012).
- [43] K. C. Cannon, *Class. Quant. Grav.* **25**, 105024 (2008).
- [44] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji, *et al.*, *Class.Quant.Grav.* **25**, 184004 (2008), arXiv:0804.0800.
- [45] B. Allen, *Phys.Rev.* **D71**, 062001 (2005), arXiv:gr-qc/0405045.
- [46] C. Hanna, *Searching for Gravitational Waves from Binary Systems in Non-Stationary Data*, Ph.D. thesis, Louisiana State University (2008).
- [47] B. Allen, J. D. Creighton, E. E. Flanagan, and J. D. Romano, *Phys. Rev.* **D67**, 122002 (2003), arXiv:gr-qc/0205015.
- [48] T. B. Littenberg and N. J. Cornish, *Phys. Rev.* **D82**, 103007 (2010), arXiv:1008.1577.
- [49] A. C. Searle, (2008), arXiv:0804.1161.